arXiv:1703.06581v1 [math.OC] 20 Mar 2017

Disaggregated Benders decomposition and lazy constraints for solving the budget-constrained dynamic uncapacitated facility location and network design problem

Robin H. Pearce & Michael Forbes

School of Mathematics and Physics, University of Queensland, Australia

Abstract

We present an approach for solving to optimality the budget-constrained Dynamic Uncapacitated Facility Location and Network Design problem (DUFLNDP). This is a problem where a network must be constructed or expanded and facilities placed in the network, subject to a budget, in order to satisfy a number of demands. With the demands satisfied, the objective is to minimise the running cost of the network and the cost of moving demands to facilities. The problem can be disaggregated over two different sets simultaneously, leading to many smaller models which can be solved more easily. Using disaggregated Benders decomposition and lazy constraints, we solve many instances to optimality that have not previously been solved. We use an analytic procedure to generate Benders optimality cuts which are provably Pareto-optimal.

Keywords: Network design, Benders decomposition, Lazy constraints, Pareto-optimality

1. Introduction

In this paper we apply Benders decomposition to a facility location and network design problem, specifically looking at a number of ways of improving convergence of the algorithm. In particular, we disaggregate the Benders sub-problems, use an alternative to the standard Benders feasibility cuts and analytically construct Benders optimality cuts. We also prove the Pareto-optimality of the analytic Benders cuts and discuss the importance of using Pareto-optimal cuts.

Facility location problems are important in many areas of both industry and government. From deciding the location of stores and warehouses, to important services such as police, fire and health, facility location problems can have a large impact on a population. Equally important are network design problems, such as vehicle routing, or utility network optimisation. While these problems

have individually been extensively researched and expanded, the combination of facility location and network design has received less attention.

The facility location problem dates back to the start of the 20th century [1], and is the basis of many more detailed problems. Benders decomposition [2] is an ideal technique for solving facility location problems, particularly the uncapacitated facility location (UFL) problem [3]. Geoffrion and Graves [4] apply Benders Decomposition to a multicommodity variant of the facility location problem to great effect. Magnanti and Wong [5] explore regular and disaggregated Benders Decomposition, apply it to the UFL problem, and propose an interior point method for accelerating convergence of the algorithm.

More recently, an efficient implementation of Benders Decomposition for the UFL is demonstrated by Fischetti, Ljubic and Sinnl (2016)[3]. They apply disaggregated Benders decomposition with a number of additional features which are useful, particularly for the UFL. Tang, Jiang and Saharidis [6] use disaggregated Benders decomposition to solve a capacitated facility location problem where the capacities could be modified for a cost. They also consider adding extra constraints to enforce feasibility and tighten the lower bound on the objective value, which are important in the application of Benders Decomposition.

Network flow and design problems have also been a major area of study over the last century. Today, many efficient methods for finding the maximum flow through a network exist [7, 8]. As such, more recent studies tend to focus on network design problems, where the network itself is optimised to achieve some goal, such as maximising the throughput of the network over time. Many of these problems are excellent candidates for Benders decomposition. Nurre, Cavdaroglu and Wallace [9] consider a problem where a utilities network has been partially destroyed, and the reconstruction must be scheduled to maximise total throughput of the network over time. Boland et. al. [10] find the optimal maintenance schedule of a network, also to maximise throughput.

Another goal may be to minimise the running cost of a network, which is often coupled with a facility- or hub-location aspect. One example of this is the Hub Line Location problem, considered by de Sá et. al. [11], where hub facilities must be built in a public transit network and connected in a line. The objective is to minimise the weighted travel time of all demands through the network. Another example is the Uncapacitated Multiple Allocation Hub Location problem considered by Camargo, Miranda Jr. and Luna [12], where hubs must be built so demands can be routed between locations via hubs. de Sá, de Camargo and de Miranda [13] apply Benders decomposition to another

hub location problem, with a number of improvements such as a "warm start", disaggregation of the sub-problems and modified feasibility cuts. Both studies apply Benders decomposition to their problems.

We are considering the budget-constrained Dynamic Uncapacitated Facility Location and Network Design Problem (DUFLNDP) presented by Ghaderi and Jabalameli [14]. The government sets a fixed budget every year for the construction of new health clinics and roads, and one must work within that budget to minimise the running cost of the network while satisfying all demand for health services by routing demand through the network to health clinics.

The remainder of this paper is structured as follows: Section 2 contains our reformulation of the DUFLNDP, which is the base model to which we apply disaggregated Benders decomposition in Section 3. This section also covers many details around the implementation of disaggregated Benders decomposition such as Pareto-optimality of Benders optimality cuts and feasibility of sub-problems. In Section 4 we describe the use of a warm start with Benders decomposition to improve the initial LP-bound. Our computational results are in Section 5, followed by a discussion of disaggregation in Section 6, before concluding with Section 7.

2. Model Formulation

Ghaderi and Jabalameli [14] introduce the budget-constrained Dynamic Uncapacitated Facility Location and Network Design problem (DUFLNDP), which is defined on a network of locations. Every location is a client, and all have the potential to host a facility for servicing clients. There is a set of potential links between locations, on which arcs of the network can be constructed.

The problem covers a number of time periods. At each time there are budgets for opening new facilities and links. Open facilities and links also have associated maintenance or operating costs, which, together with the demand routing costs, form the total cost which is to be minimised.

The main assumptions in this problem are:

- Facilities and links have unlimited capacity
- Once opened, facilities and links will remain open until at least the end of the planning horizon
- Facilities and links are opened instantaneously between time periods

Our notation is slightly different from Ghaderi and Jabalameli [14], in particular the variable names. We also present a simplified version of the budgetary constraints which achieve the same outcome. The time periods we are optimising over start at 1, and if a network exists already, we denote that as being at time 0. We now present the model formulation:

Sets

N	Set of network nodes. These include clients and facilities
A	Set of network arcs, both existing and potential. $A \subseteq N \times N$
T	Set of time periods

Parameters

d_{kt}	Demand of client $k \in N$ at time $t \in T$
g_{it}	Cost of opening facility at node $i \in N$ at time $t \in T$
c_{ijt}	Cost of constructing arc $(i,j) \in A$ at time $t \in T$
ρ_{ijt}	Cost per unit of routing demand on arc $(i,j) \in A$ at time $t \in T$
f_{it}	Operating cost of open facility $i \in N$ at time $t \in T$
h_{ijt}	Operating cost of open arc $(i,j) \in A$ at time $t \in T$
$\bar{B}_t$	Available budget for opening facilities at time $t \in T$
$\hat{B}_t$	Available budget for opening arcs at time $t \in T$

Variables

W_{it}	1 if facility $i \in N$ is open at time $t \in T$, 0 otherwise
X_{ijt}	1 if arc $(i,j) \in A$ is open at time $t \in T$, 0 otherwise
Z_{ijkt}	Fraction of demand of client $k \in N$ travelling along arc $(i,j) \in A$ at time $t \in T$
U_{it}	1 if facility $i \in N$ is constructed at time $t \in T$, 0 otherwise
V_{ijt}	1 if arc $(i,j) \in A$ is constructed at time $t \in T$, 0 otherwise

Objective

$$\text{Minimise} \sum_{t \in T} \left(\sum_{i \in N} f_{it} W_{it} + \sum_{k \in N} \sum_{(i,j) \in A} \rho_{ijt} d_{kt} Z_{ijkt} + \sum_{\substack{(i,j) \in A \\ i<j}} h_{ijt} X_{ijt} \right) \tag{1}$$

Constraints

$$W_{kt} + \sum_{j \in N} Z_{kjkt} \geq 1 \qquad \forall k \in N, \forall t \in T \tag{2}$$

$$\sum_{j \in N} Z_{jikt} \leq \sum_{j \in N} Z_{ijkt} + W_{it} \qquad \forall i, k \in N, i \neq k, \forall t \in T \tag{3}$$

$$Z_{jkkt} = 0 \qquad \forall k \in N, \forall j \in N, \forall t \in T \tag{4}$$

$$Z_{ijkt} \leq X_{ijt} \qquad \forall (i,j) \in A, \forall k \in N, \forall t \in T \tag{5}$$

$$W_{i,t-1} + U_{it} = W_{it} \qquad \forall i \in N, \forall t \in T \tag{6}$$

$$X_{ij,t-1} + V_{ijt} = X_{ijt} \qquad \forall (i,j) \in A, \forall t \in T \tag{7}$$

$$\sum_{t'=1}^{t} \sum_{i \in N} g_{it'} U_{it'} \leq \sum_{t'=1}^{t} \bar{B}_{t'} \qquad \forall t \in T \tag{8}$$

$$\sum_{t'=1}^{t} \sum_{(i,j) \in A} c_{ijt'} V_{ijt'} \leq \sum_{t'=1}^{t} \hat{B}_{t'} \qquad \forall t \in T \tag{9}$$

$$X_{ijt} = X_{jit} \qquad \forall (i,j) \in A, i < j, \forall t \in T \tag{10}$$

$$W_{it} \in \{0,1\}, U_{it} \in \{0,1\} \qquad \forall i \in N, \forall t \in T \tag{11}$$

$$X_{ijt} \in \{0,1\}, V_{ijt} \in \{0,1\}, Z_{ijkt} \geq 0 \qquad \forall (i,j) \in A, \forall t \in T, \forall k \in N \tag{12}$$

The objective function (1) is the sum of three costs: the facility operating costs, the cost of routing demand to other facilities and the arc operating costs. Constraints (2) say that if a node k has an open facility, then it services its own demand. If not, all demand must leave the node. Constraints (3) are flow-conservation constraints at the nodes. Constraints (4) ensure demand can not be returned to the node of origin, thus eliminating cycles. Constraints (5) restrict the routing of demand to open arcs only. Constraints (6) and (7) control the opening of facilities and arcs based on the relevant construction variables, and constraints (8) and (9) ensure that the budget is not exceeded in any time period. Finally, constraints (10) enforce bi-directionality of the arcs.

3. Disaggregation and Benders Decomposition

In this problem, the variables Z_{ijkt} are continuous, where all others are integer (binary). The constraints which contain the continuous variables are (2-5), and these constraints are separate for

each $k \in N$ and $t \in T$. Thus it is possible to disaggregate the sub-problems by time and facility. A discussion of disaggregation level can be found in Section 6.1. The goal of each sub-problem is to find the cheapest way of servicing the demand of that facility at that time. There are two possibilities for this: either the site is a facility and can service its own demand for free, or the demand is routed to the nearest (cheapest) open facility.

3.1. Benders Master Problem

We denote the contribution of the sub-problem (k, t) as θ_{kt}. The master problem is:

$$\text{Minimise} \sum_{t \in T} \left(\sum_{i \in N} f_{it} W_{it} + \sum_{k \in N} d_{kt} \theta_{kt} + \sum_{\substack{(i,j) \in A \\ i<j}} h_{ijt} X_{ijt} \right) \tag{MP-OBJ}$$

Subject to:

$$W_{i,t-1} + U_{it} = W_{it} \qquad \forall i \in N, \forall t \in T \tag{M1}$$

$$X_{ij,t-1} + V_{ijt} = X_{ijt} \qquad \forall (i,j) \in A, \forall t \in T \tag{M2}$$

$$\sum_{t'=1}^{t} \sum_{i \in N} g_{it'} U_{it'} \leq \sum_{t'=1}^{t} \bar{B}_t \qquad \forall t \in T \tag{M3}$$

$$\sum_{t'=1}^{t} \sum_{(i,j) \in A} c_{ijt'} V_{ijt'} \leq \sum_{t'=1}^{t} \hat{B}_t \qquad \forall t \in T \tag{M4}$$

$$\theta_{kt} \geq \text{BendersOptimalityCut}(m, \mathbf{W}, \mathbf{X}, k, t) \qquad \forall k \in N, \forall t \in T, \forall m \in \{1, ..., M\} \tag{BOC}$$

$$\text{BendersFeasibilityCut}(p, \mathbf{W}, \mathbf{X}) \qquad \forall p \in \{1, ..., P\} \tag{BFC}$$

$$W_{it} \in \{0,1\}, U_{i,t} \in \{0,1\}, \theta_{it} \geq 0 \qquad \forall i \in N, \forall t \in T \tag{M5}$$

$$X_{ijt} \in \{0,1\}, V_{ijt} \in \{0,1\} \qquad \forall (i,j) \in A, \forall t \in T \tag{M6}$$

Constraints (BOC) represent the disaggregated Benders cuts, which are added as necessary after solving the associated sub-problems, which will be covered in Section 3.4. Similarly, constraints (BFC) represent the added constraints required for feasible sub-problems. M and P are the number of added Benders optimality and feasibility cuts respectively. For a feasible integer solution, W^* and X^*, we solve each of the sub-problems and calculate their dual variables. If necessary, we add more Benders optimality or feasibility cuts.

3.2. Initial feasibility

For the solution to be feasible, it must be possible to service the demand of every client for every time period. With the current master problem and Benders optimality cuts, it is possible for the master problem to be solved to optimality while one or more sub-problems are infeasible. The standard way of overcoming this is to add Benders feasibility cuts, however these are often ineffective [13].

A second option is to modify the master problem to ensure the sub-problems will always be feasible. Since links and facilities are only constructed, never destroyed, if the network is feasible in the first time period, it will be feasible for every time period. To ensure this happens, we modify the model to make the first time period a special case. The objective, parameters and variables remain unchanged, we only modify some constraints and add new ones. The modified and new constraints are:

Constraints

$$Z_{ijkt} \leq X_{ijt} \qquad \forall (i,j) \in A, \forall k \in N, \forall t \in T \tag{5a}$$

$$X_{ij,t-1} + V_{ijt} = X_{ijt} \qquad \forall (i,j) \in A, i < j, \forall t \in T, t > 2 \tag{7a}$$

$$X_{ijt} = X_{jit} \qquad \forall (i,j) \in A, i < j, \forall t \in T, t > 1 \tag{10a}$$

$$X_{ij,1} + X_{ji,1} \leq 1 \qquad \forall (i,j) \in A, i < j \tag{13}$$

$$X_{ij,0} + V_{ij,1} = X_{ij,1} + X_{ji,1} \qquad \forall (i,j) \in A, i < j \tag{14}$$

$$X_{ij,1} + X_{ji,1} + V_{ij,2} = X_{ij,2} \qquad \forall (i,j) \in A, i < j \tag{15}$$

$$\sum_{\substack{j \in N \\ (i,j) \in A}} X_{ij,1} + W_{i,1} \geq 1 \qquad \forall i \in N \tag{16}$$

$$\sum_{i \in N} W_{i,1} \geq 1 \tag{17}$$

The modification to constraint (5) enforces directionality of arcs in the first time period, (7,10) are modified appropriately, and the addition of (13) ensures only one direction is allowed for that time. Constraints (14-15) handle the budget constraints, to ensure that if a direction is built in the first time period, the opposite direction will be built for free in the second time period.

These modifications allow us to add constraints (16-17), which ensure that each location has either a facility at the location or an arc leaving the location, and that at least one facility must exist, respectively. This way, either a node is a facility, or it is connected to a node which is either a facility, or connected to a node... and so on. This only fails if a cycle occurs where multiple nodes are connected to each other and none have facilities, so cycle-breaking may be necessary. This change in formulation is more useful in the instances when there is no pre-existing network, as when there are fixed elements of the network there is less choice in its design.

3.3. IIS feasibility cuts

To handle the case where cycles occur, we add cycle-breaking feasibility cuts, where the sum of facilities in the cycle plus the sum of arcs leaving the cycle must be at least one. Using Gurobi, we compute the Irreducible Inconsistent Subsystem (IIS), which is "a subset of the constraints and variable bounds of the original model. If all constraints in the model except those in the IIS are removed, the model is still infeasible. However, further removing any one member of the IIS produces a feasible result."[15]

The IIS is then a collection of capacity constraints on nodes and arcs which, when lifted, make the sub-problem feasible. This leaves us with the nodes and arcs which can be expanded or added to resolve the infeasibility. We then add a feasibility cut of the form:

$$\sum_{i \in \text{IIS}} W_{i0} + \sum_{(a,b) \in \text{IIS}} X_{ab0} \geq 1 \tag{18}$$

This ensures that enough facilities and arcs will be opened that the demand from the infeasible source node can be served.

3.4. Sub-Problems

If we have a feasible solution for the integer variables W^* and X^*, we can solve the sub-problems as a collection of linear programs. Since d_{kt} only depends on k and t, we can leave it out of the objective of the sub-problem and instead apply it to the objective of the master problem. The contribution of each sub-problem to the master problem is represented by θ_{kt}. For each $k \in N$ and $t \in T$ we have the sub-problem:

$$\text{Minimise} \sum_{(i,j) \in A} \rho_{ijt} Z_{ijkt} \tag{SP-OBJ}$$

Subject to:

$$-\sum_{j\in N} Z_{kjkt} \leq W^*_{kt} - 1 \tag{S1}$$

$$\sum_{j\in N} Z_{jikt} - \sum_{j\in N} Z_{ijkt} \leq W^*_{it} \qquad \forall i \in N \setminus \{k\} \tag{S2}$$

$$Z_{jkkt} = 0 \qquad \forall j \in N \tag{S3}$$

$$Z_{ijkt} \leq X^*_{ijt} \qquad \forall (i,j) \in A \tag{S4}$$

$$Z_{ijkt} \geq 0 \qquad \forall (i,j) \in A \tag{S5}$$

Constraint (S1) has been rearranged to show its similarity to (S2), which we will take advantage of when formulating the explicit dual. There are two collections of dual variables that we are interested in: γ_i for each node constraint (S1,S2) and λ_{ij} for each arc constraint (S4). These variables then lead to the following dual formulation:

$$\text{Maximise } \gamma_k - \sum_{i\in N} \gamma_i W^*_{it} - \sum_{(i,j)\in A} \lambda_{ij} X^*_{ijt} \tag{D-OBJ}$$

Subject to:

$$\rho_{ijt} + \lambda_{ij} + \gamma_j - \gamma_i \geq 0 \qquad \forall (i,j) \in A, j \neq k \tag{D1}$$

$$\lambda_{ij} \geq 0, \gamma_i \geq 0 \qquad \forall (i,j) \in A, \forall i \in N \tag{D2}$$

The reason that constraints (D1) do not apply when $j = k$ is because in those cases γ_j is replaced by the unbounded dual variable associated with constraints (S3), and since this dual variable does not appear in the objective function, it can be set to positive infinity, thus ensuring feasibility of the dual constraints for those arcs.

Magnanti and Wong [5] describe a natural interpretation of the dual variables for the UFL. In a similar way, the dual variables of this problem have a natural interpretation. Here each γ_i represents the saving associated with opening a facility at location i, and λ_{ij} is the saving from opening a new arc between i and j. Constraint (D1) ensures that the reduced cost of each arc is non-negative. This formulation yields the following Benders cut:

$$\theta_{kt} \geq \gamma_k - \sum_{i\in N} \gamma_i W_{it} - \sum_{(i,j)\in A} \lambda_{ij} X_{ijt} \tag{BC}$$

One can solve these sub-problems as linear programs and extract the dual variables provided by the solver in order to construct a Benders cut. Alternatively, one can solve the sub-problems and produce the required dual variables analytically.

3.5. Analytic solution to Benders sub-problems

Each sub-problem is a shortest path problem. For each location, one must find the cheapest way of servicing its demand, either at a facility at the source or by routing the demand to another location with a facility. Each sub-problem is indexed by k and t, where k is the source node and t is the time period.

Magnanti and Wong [5] note that the analytic dual variables for the UFL problem have a natural interpretation; the dual variables for the DUFLNDP also have a natural interpretation. γ_i represents the saving associated with opening a facility at location i and λ_{ij} the saving from opening an arc from i to j. Magnanti and Wong also demonstrate for the UFL problem that there are two different possible analytic Benders cuts which can be used. The DUFLNDP shares this property, as we now show.

3.5.1. First analytic Benders cut

The algorithm for computing the dual variables can be found in Algorithm 1. Note that it assumes that the facility at the source node k is closed, as otherwise the solution is trivial. We begin by constructing a shortest path tree from the source location, giving each node a distance D_i from the source node. If there is no path between k and i, then $D_i = \infty$. Now, the value of γ_k is assigned the length of the shortest path to the nearest open facility i^*, that is, $\gamma_k = \min\{D_i | W^*_{it} = 1\} \equiv D_{i^*}$. For all other nodes, $\gamma_i = \max(0, \gamma_k - D_i)$.

Next we calculate the values for the dual variables λ_{ij}, associated with the arcs $(i, j) \in A$. For all open arcs ($X^*_{ijt} = 1$), $\lambda_{ij} = 0$. For any closed arc (i, j), $\lambda_{ij} = \max(0, \gamma_i - \gamma_j - \rho_{ij})$.

Theorem 1. *The dual variables calculated using Algorithm 1 are dual optimal.*

Proof. For these dual variables to form a dual feasible solution, they must satisfy the Constraints (D1). The constraints are trivially satisfied for any arc where $\gamma_i = 0$. For all closed arcs, $\lambda_{ij} \geq \gamma_i - \gamma_j - \rho_{ijt}$, which satisfies Constraint (D1).

For any open arc (i, j), $\lambda_{ij} = 0$, so we must show that $\rho_{ij} + \gamma_j - \gamma_i \geq 0$. By the property of the shortest path distances, $D_j \leq D_i + \rho_{ij}$, or $\rho_{ij} + D_i - D_j \geq 0$. We also have, by construction of the

Algorithm 1 Algorithm for computing dual variables for first analytic Benders optimality cut for sub-problem (k, t), assuming the sub-problem is feasible.

Begin with master problem solution $W^*_{it}, X^*_{abt}, \theta^*_{kt}\ \forall i \in N, \forall (a,b) \in A$

Compute shortest distance D_i from k to i for all nodes $i \in N \setminus \{k\}$

$\gamma_k = \min_j \{D_j | W^*_{jt} = 1, j \in N \setminus \{k\}\}$

for $i \in N$ **do**

  $\gamma_i \leftarrow \max(0, \gamma_k - D_i)$

for $(i,j) \in A$ **do**

  if $X^*_{ijt} = 0$ **then**

    $\lambda_{ij} \leftarrow \max(0, \gamma_i - \gamma_j - \rho_{ijt})$

  else

    $\lambda_{ij} \leftarrow 0$

Add Constraint $\theta_{kt} \geq \bar{\theta}_{kt} - \sum_{i \in N} \gamma_i W_{it} - \sum_{(i,j) \in A} \lambda_{ijt} X_{ijt}$

dual variables, that $\gamma_j \geq \gamma_k - D_j$, or $D_j \geq \gamma_k - \gamma_j$. Finally, we are only considering where $\gamma_i > 0$, and in this case we have that $\gamma_i = \gamma_k - D_i$. Combining these, we get:

$$\begin{aligned} \rho_{ij} + \gamma_j - \gamma_i &= \rho_{ij} + \gamma_j - (\gamma_k - D_i) \\ &\geq \rho_{ij} + (\gamma_k - D_j) - (\gamma_k - D_i) \\ &\geq \rho_{ij} + D_i - D_j \\ &\geq 0 \end{aligned}$$

So the dual variables obtained from Algorithm 1 are dual feasible. The objective value given by these dual variables is the same as the optimal objective value of the primal problem, which is the length of the shortest path ($D_{i^*} = \gamma_k$), and for all nodes, either $\gamma_i = 0$ or $W^*_{it} = 0$, and likewise for arcs. Thus the dual variables form a dual optimal solution, and may be used to add a Benders optimality cut to the master problem. □

The cuts generated using these dual variables are Pareto-optimal, which is important for improving convergence of the master problem [5]. We prove they are Pareto-optimal in Section 3.6. In terms of the natural interpretation, we place much of the savings on the arcs, since if facilities beyond closed arcs are opened, no saving will be obtained until the arcs are open. As such, the first

Algorithm 2 Algorithm for computing dual variables for second analytic Benders optimality cut for sub-problem (k, t), assuming the sub-problem is feasible.

Begin with master problem solution $W^*_{it}, X^*_{abt}, \theta^*_{kt}\ \forall i \in N, \forall (a, b) \in A$

Compute shortest distance $\bar{D}_i$ from i to an open facility for all nodes $i \in N$

Open all arcs $(i, j) \in A$

Compute shortest distance D_i from k to i for all nodes $i \in N \setminus \{k\}$

$\gamma_k = \bar{D}_k$

for $i \in N$ **do**

    $\gamma_i \leftarrow \max(0, \min(\gamma_k - D_i, \bar{D}_i))$

for $(i, j) \in A$ **do**

    if $X^*_{ijt} = 0$ **then**

        $\lambda_{ij} \leftarrow \max(0, \gamma_i - \gamma_j - \rho_{ijt})$

    else

        $\lambda_{ij} \leftarrow 0$

Add Constraint $\theta_{kt} \geq \bar{\theta}_{kt} - \sum_{i \in N} \gamma_i W_{it} - \sum_{(i,j) \in A} \lambda_{ijt} X_{ijt}$

analytic Benders cut is referred to as a λ-heavy cut. While this Benders cut is Pareto-optimal, it may not be the only Pareto-optimal analytic cut which can be derived from a given solution.

3.5.2. Second analytic Benders cut

Another algorithm for computing dual variables for an analytic Benders optimality cut is Algorithm 2. The construction of the second analytic cut is similar to the first, except here we also calculate the distance from each node to their closest open facility, $\bar{D}_i$. Then, before computing the distance of all nodes from the source k, we open all arcs. Again, γ_k is the length of the shortest path, this time given by $\bar{D}_k$.

Now, $\gamma_i = \max(0, \min(\gamma_k - D_i, \bar{D}_i))$, that is, it is the smaller of either the saving from opening a facility at node i assuming all arcs are open, or the shortest distance to an open facility under the current configuration of arcs. For nodes further away than i^*, $\gamma_i = 0$ again. The values of λ_{ij} are calculated in the same way as before. These dual variables also form a dual feasible solution for similar reasons as the first.

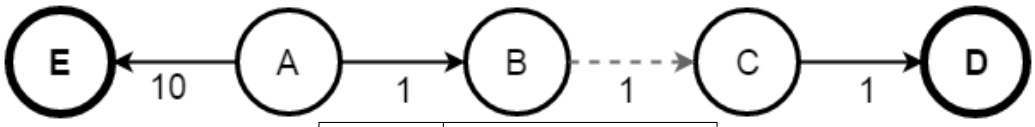

	γ_A	γ_B	γ_C	λ_{BC}
λ-heavy	10	9	0	8
γ-heavy	10	9	1	7

Figure 1: Example network with dual variables from both analytic cuts. Node A is the source node. Nodes D and E are open facilities, and the arc between B and C is currently closed.

Theorem 2. *The dual variables computed using Algorithm 2 are dual optimal*

Proof. Once again, the constraints are satisfied trivially for $\gamma_i = 0$ and by construction for closed arcs. For open arcs, $\lambda_{ij} = 0$, so again we need to prove that $\rho_{ij} + \gamma_j - \gamma_i \geq 0$. Also by the properties of shortest distances, we have $\bar{D}_i \leq \bar{D}_j + \rho_{ij}$, or $\rho_{ij} + \bar{D}_j - \bar{D}_i \geq 0$. For these dual variables, when $\gamma_i > 0$, we have that $\gamma_i \leq \gamma_k - D_i$ and $\gamma_i \leq \bar{D}_i$. Finally, either $\gamma_j \geq \gamma_k - D_j$ or $\gamma_j \geq \bar{D}_j$.

If $\gamma_j \geq \gamma_k - D_j$, the proof for the first set of dual variables holds since $\gamma_i \leq \gamma_k - D_i$. When $\gamma_j \geq \bar{D}_j$, the fact that $\gamma_i \leq \bar{D}_i$ leads to:

$$\begin{aligned} \rho_{ij} + \gamma_j - \gamma_i &\geq \rho_{ij} + \bar{D}_j - \gamma_i \\ &\geq \rho_{ij} + \bar{D}_j - \bar{D}_i \\ &\geq 0 \end{aligned}$$

So the dual variables satisfy the dual constraints for all arcs. For the same reasons as before, these dual variables give the primal objective value for the current master problem solution, making them a dual optimal solution. □

This analytic solution also yields a Pareto-optimal Benders cut, which can be proven in a similar way to the first analytic solution. Where the first was a λ-heavy cut, the second is called a γ-heavy cut, since we focus on the saving associated with opening facilities assuming the arcs are already open. We now have (up to) two unique Benders cuts that can be added at each integer solution for each node k and time t. If the two cuts are different, then they will hold exactly for different, overlapping sets of master problem solutions.

Consider the network in Figure 1. Both cuts will be tight for the cases where (B,C) and C are the same, either both open or both closed. However, only the λ-heavy cut will be tight when C

is open but (B,C) is closed, and only the γ-heavy cut will be tight when (B,C) is open but C is closed. As such, it can be beneficial to add both cuts if possible, as can be seen in the Results section. Note that these Benders cuts are not guaranteed to be different for all solutions, and it is simple to construct a network where both algorithms yield the same Benders cut. As such, one should always take care not to add duplicate Benders cuts to their model, to avoid unnecessarily burdening the solver.

3.6. Pareto-optimality of the analytically-derived Benders cuts

Magnanti and Wong [5] describe the importance of using Pareto-optimal cuts when using Benders decomposition. In this section we show that the analytic Benders cuts are Pareto-optimal. Since a Benders cut is a linear function of the current network configuration, it can be described as $\theta \geq \bar{\theta}(y)$, for $y \in Y$ where Y is the set of all feasible solutions to the master problem. Let the contribution to the objective value for network configuration y be given by $\bar{\theta}^*(y)$. The following definitions are paraphrased from Magnanti and Wong:

Definition 1. *A Benders cut $\theta \geq \bar{\theta}^a(y)$ dominates another Benders cut $\theta \geq \bar{\theta}^b(y)$ if $\bar{\theta}^a(y) \geq \bar{\theta}^b(y)$ for all feasible $y \in Y$ and is a strict inequality for at least one feasible y.*

This definition leads to the following lemma:

Lemma 1. *If $\theta \geq \bar{\theta}^a(y)$ is dominated by $\theta \geq \bar{\theta}^b(y)$, then for all feasible solutions y^i where $\bar{\theta}^a(y^i) = \bar{\theta}^*(y^i)$, $\bar{\theta}^b(y^i) = \bar{\theta}^*(y^i)$.*

This is easy to see, since $\bar{\theta}^b(y^i) \leq \bar{\theta}^*(y^i)$ by definition of being a valid Benders cut, and $\bar{\theta}^b(y^i) \geq \bar{\theta}^a(y^i) = \bar{\theta}^*(y^i)$ by definition of being a dominating cut.

Definition 2. *A Benders cut $\theta \geq \bar{\theta}^a(y)$ is Pareto-optimal if it is not dominated by any other Benders cuts.*

One can prove that a Benders cut is Pareto-optimal by assuming that there exists another cut which dominates it, finding enough points $y \in Y$ where the Pareto-optimal cut equals the objective value, and specifying that the dominating cut must also equal the objective value at these points. This leads to all terms of the dominating cut being fixed to those of the Pareto-optimal cut. Thus there are no cuts which dominate the original cut, and it is Pareto-optimal. We now show that our Benders optimality cuts are Pareto-optimal.

Theorem 3. *Benders optimality cuts derived using Algorithm 1 are Pareto-optimal*

Proof. All Benders cuts for this problem are of the form:

$$\theta_{kt} \geq \gamma_k - \sum_{i \in N} \gamma_i W_{it} - \sum_{(i,j) \in A} \lambda_{ij} X_{ijt} \tag{19}$$

Let the dual variables associated with the first analytic Benders cut (i.e. the cut generated by Algorithm 1) be $\bar{\gamma}_i$ and $\bar{\lambda}_{ij}$. In the current solution, the closest open facility is i^* ($d_{i^*} = \bar{\gamma}_k$).

We assume that the current cheapest facility i^* is not the same as the location, i.e. $i^* \neq k$. If $i^* = k$, then the cut is trivial and not Pareto-optimal. We also assume that there exists another location $j \in N$ such that $d_j < d_{i^*}$, since otherwise the cut will again be trivial and not Pareto-optimal. We begin by defining some partitions of the nodes and arcs of the problem:

$\mathbf{F}^o = \{i | W^*_{it} = 1\}$, the set of open facilities,

$\mathbf{F}^c = \{i | W^*_{it} = 0\}$, the set of closed facilities,

$\mathbf{F}^+ = \{i | d_i \geq \bar{\gamma_k}\}$, the set of facilities at equal or greater distance than i^*,

$\mathbf{F}^- = \{i | d_i < \bar{\gamma_k}\}$, the set of facilities closer than i^*,

$\mathbf{L}^o = \{i | X^*_{ijt} = 1\}$, the set of open links, and

$\mathbf{L}^c = \{i | X^*_{ijt} = 0\}$, the set of closed links

Now assume there exists a Benders cut using the dual variables $\hat{\gamma}_i$ and $\hat{\lambda}_{ij}$, which dominates the first analytic Benders cut. As they are both Benders cuts, they must both equal the objective value, and thus each other, for the current solution to the master problem. If we open a facility at the source, k, the objective value will be zero and the first analytic cut will be tight, so the dominating cut must also equal zero for this solution. This leads to:

$$\begin{aligned} 0 =& \bar{\gamma_k} - \bar{\gamma_k} - \sum_{\substack{i \in N \\ i \neq k}} \bar{\gamma}_i W^*_{it} - \sum_{(i,j) \in A} \bar{\lambda}_{ij} X^*_{ijt} \\ =& 0 - \sum_{\substack{i \in \mathbf{F}^o \\ i \neq k}} \bar{\gamma}_i - \sum_{(i,j) \in \mathbf{L}^o} \bar{\lambda}_{ij} X^*_{ijt} \end{aligned}$$

Since $\bar{\gamma}_i \geq 0$ and $\bar{\lambda}_{ij} \geq 0$, we have that $\bar{\gamma}_i = \hat{\gamma}_i = 0 \ \forall i \in \mathbf{F}^o$ and $\bar{\lambda}_{ij} = \hat{\lambda}_{ij} = 0 \ \forall (i,j) \in \mathbf{L}^o$. Note that $i^* \in \mathbf{F}^o$, and so $\gamma_{i^*} = 0$. Returning to the current solution, the two cuts must equal each other, and for both cuts, either $\gamma_i = 0$ or $W_{it} = 0 \ \forall i \in N$, and similarly for arcs, so we have that $\hat{\gamma}_k = \bar{\gamma}_k$.

Now, for any other location $i \in N$, if we open a facility at i, the first analytic cut will still be tight, so the dominating cut must also be tight. Since the only changes in both cuts is W_{it}, we have that $\bar{\gamma}_i = \hat{\gamma}_i \ \forall i \in N$. All that remains is to show that $\bar{\lambda}_{ij} = \hat{\lambda}_{ij} \ \forall (i,j) \in \mathbf{L}^c$.

For any arc $(i,j) \in \mathbf{L}^c$ where $\hat{\gamma}_i \leq \rho_{ijt}$, $\hat{\lambda}_{ij} = 0$, which will be tight since even if a facility were opened at j, it would still be further away than the closest open facility, so $\bar{\lambda}_{ij}$ will also be zero. The other case where $\hat{\lambda}_{ij} = 0$ is when $\hat{\gamma}_j > \hat{\gamma}_i - \rho_{ijt}$, that is, the arc does not create a short-cut in the network. In this case, opening the arc does not change the objective value and the first analytic cut will be tight, so $\bar{\lambda}_{ij} = 0$ for all arcs where $\hat{\lambda}_{ij} = 0$.

If $\hat{\lambda}_{ij} > 0$, then $\hat{\lambda}_{ij} = \hat{\gamma}_i - \hat{\gamma}_j - \rho_{ijt}$ and $\hat{\gamma}_i > \rho_{ijt}$. If we open the arc (i,j) and the facility j, then the first analytic cut will be:

$$\begin{aligned}\hat{\gamma}_k - \hat{\gamma}_j - \hat{\lambda}_{ij} =& \hat{\gamma}_k - \hat{\gamma}_j - (\hat{\gamma}_i - \hat{\gamma}_j - \rho_{ijt}) \\ =& \hat{\gamma}_k - \hat{\gamma}_i + \rho_{ijt}\end{aligned}$$

which is the length of the shortest path between k and i plus the length of the arc from i to j. Since $\hat{\gamma}_i > \rho_{ijt}$, this will be lower than the original path length, and thus j will be closer than i^* to k. So the first analytic cut is tight at these points, and the dominating cut must also be tight. Since $\bar{\gamma}_j = \hat{\gamma}_j$ for all $j \in N$, we have:

$$\begin{aligned}\hat{\gamma}_k - \hat{\gamma}_j - \hat{\lambda}_{ij} =& \bar{\gamma}_k - \bar{\gamma}_j - \bar{\lambda}_{ij} \\ =& \hat{\gamma}_k - \hat{\gamma}_j - \bar{\lambda}_{ij} \\ \hat{\lambda}_{ij} =& \bar{\lambda}_{ij}\end{aligned}$$

So $\bar{\lambda}_{ij} = \hat{\lambda}_{ij} \ \forall (i,j) \in A$, and thus the dominating cut is identical to the first analytic cut. So there are no cuts which dominate the first analytic cut, and thus it is Pareto-optimal. □

The same method can be used to prove that the second analytic Benders cut is also Pareto-optimal by selecting different points.

3.7. Budget cover inequalities

In addition to this, we also add inequalities on the budget variables, U and V, to potentially tighten the relaxed problem. The budget constraints are effectively a knapsack problem, and as such we can add cover inequalities similar to those described by Gu, Nemhauser and Savelsbergh

[16]. After a solution to the relaxed problem is found, we check, for each time period, which facilities and arcs have been partially or wholly constructed. We sum the variables over all facilities/links and time periods up to and including the current time period, and if this is not an integer value, then some facilities or links have been partially opened.

S is the sum of facilities/links that have been opened up to this point in time. We then order the facilities/links from cheapest to most expensive to open, and if the sum of opening costs of the first $\lceil S \rceil$ facilities/links is greater than the available budget, we add a new constraint of the form:

$$\sum_{t'=1}^{t} \sum_{i \in \bar{S}} U_{it'} \leq \lfloor S \rfloor \qquad (20) \qquad\qquad \sum_{t'=1}^{t} \sum_{\substack{(i,j) \in \bar{S} \\ i<j}} V_{ijt'} \leq \lfloor S \rfloor \qquad (21)$$

where $\bar{S}$ is the cheapest $\lceil S \rceil$ facilities/links. If it is impossible to open all facilities/links in $\bar{S}$, then turning one off to open another facility/link which is more expensive will not be possible either. Thus, these cuts can be lifted to include all facilities/links more expensive than the most expensive member of $\bar{S}$.

4. Warm start

One potential problem from applying disaggregated Benders Decomposition to this model is that the initial LP bound is extremely loose. While it is expected that the LP bound will be lower for the Benders master problem because it is a relaxation of the original problem, for this model it is significantly lower. This can be overcome by using a "warm-start"[13, 17], which involves solving the linear relaxation of the problem and using the results to add Benders cuts to the master problem. Performing this repeatedly until the bound does not increase substantially, or no more cuts are added, significantly tightens the bound and reduces the runtime of the solver.

This yields significant improvements to the runtime of the program, however it is sometimes more useful to use continuous analogues of the Pareto-optimal analytic Benders cuts in the warm start. Because of this, we analytically construct the dual variables to be used in the pre-cuts. This

yields the strongest cuts possible, which can improve the solution speed of the master problem.

4.1. Feasibility of sub-problems

The first thing to check, as with the main algorithm, is the feasibility of the sub-problems. In the warm start, because all variables are continuous and not integer (or binary), the arcs of the network are allowed to be partially open, and likewise for facilities. As such, it is no longer enough that there be a path to an open facility, instead it must be possible to route all demand to facilities simultaneously. As with the main problem, infeasible sub-problems will only occur when a cycle exists in the network. The IIS feasibility cuts are capable of handling the relaxed problem, and as such are always used in the warm start.

4.2. Solution of the sub-problems

After having ensured the feasibility of the relaxed solutions to the master problem, we solve the flow sub-problems as LP's and extract the paths from these results. When more than one path is required, it is because partially opened arcs or facilities are restricting the flow of demand. In most cases, the longest path will not have any of these restricting factors.

If there are n paths in the solution, there will be at least $n-1$ restricting factors. These restricting factors, denoted by the set $\mathcal{C}$, correspond to potential non-zero values for γ or λ dual variables, and as such there are several constraints on these values. The first is that the sum of the dual variables corresponding to the arcs and final facility of each path must equal the saving from travelling along the path. That is, given a path p of length L_p which ends at node dest_p, and the set of arcs on that path A_p:

$$\gamma_k - L_p = \gamma_{\text{dest}_p} + \sum_{(a,b)\in A_p} \lambda_{ab} \tag{22}$$

This equation ensures the reduced cost of each path is zero. In most cases, γ_k will be the length of the longest path (opening a facility allows demand from the longest path to be serviced at the source), however in the case where the longest path has restricting factors, the RHS of the above equation will be non-zero for the longest path and γ_k will be greater than L_p.

Another condition is that the value of the Benders cut must be equal to the objective value of the sub-problem for the current master problem solution. This is necessary for the dual variables

to form a dual-optimal solution. The final condition is that the reduced cost of each arc is non-negative. If a solution is found which satisfies these three conditions, then it is dual optimal, and the dual variables can be used to construct a Benders cut.

If there are n paths and $n-1$ restricting factors, then these dual variables can be calculated directly by solving equation (22) for all paths simultaneously. In this case, the matrix will be non-singular and thus the values of the dual variables for all restricting factors can be determined. However, as there is much degeneracy in network flow problems, often there will be more than $n-1$ restricting factors for n paths. This occurs when a path has two or more restricting factors which lie only on that path. In this case, one can either determine which $n-1$ factors to use by eliminating any "extra" factors, or one can solve the following linear program:

$$\text{Minimise } \gamma_k \tag{23}$$

Subject to:

$$L_p - \gamma_k + \gamma_{\text{dest}_p} + \sum_{\substack{(a,b)\in\mathcal{C} \\ (a,b)\in A_p}} \lambda_{ab} = 0 \qquad \forall p \in P \tag{24}$$

$$\gamma_k - \sum_{i\in N} \gamma_i \bar{W}_{it} - \sum_{(i,j)\in A} \lambda_{ij} \bar{X}_{ijt} = \sum_{(i,j)\in A} \rho_{ijt} d_{kt} \bar{Z}_{ijkt} \tag{25}$$

$$\gamma_j - \gamma_i + \lambda_{ij} + \rho_{ijt} d_{kt} \geq 0 \qquad \forall (i,j) \in A \tag{26}$$

The effect of constructing analytic warm start Benders cuts this way can be seen in the Results section.

5. Results

We are comparing three different formulations using the public data set from Ghaderi and Jabalameli [14]. The tests are performed on a PC running Windows 8.1 with an Intel Core i7-3770 quad-core at 3.40GHz and 16GB of RAM. The implementations are written in Python 2.7 and use the Gurobi 6.5 [15] optimisation package. All software used is 64-bit. The maximum runtime for each instance is $50|N||T|$ seconds, where $|N|$ is the number of nodes and $|T|$ is the number of time periods, which is consistent with Ghaderi and Jabalameli [14].

All instances are grouped in threes, where each instance in a group is on the same network, tested over five, 10 and 20 time periods. Each instance has two cases: one where a network already

exists, and one where it must be created from scratch. Table 1 shows the number of nodes, links and time periods of each instance, which can be used to calculate the runtime of each instance.

We start with the straightforward MIP implementation with no improvements. We then compare it to two different implementations of disaggregated Benders decomposition: the first (DBD) is the standard implementation of disaggregated Benders decomposition with regular feasibility cuts, the second (Accelerated DBD) includes the initial time feasibility with IIS feasibility cuts. Both implementations use warm starts without budget cover constraints. Finally, we compare the different possible cuts for the warm start and callback, and evaluate the effectiveness of the budget cover constraints.

Tables 2 and 3 show the differences between the MIP, DBD and Accelerated DBD implementations on the existing and new network cases respectively. For the majority of cases, Benders decomposition is better than the straightforward MIP, especially on larger networks. There is only a minor benefit to Accelerated DBD over DBD for the existing networks, however on the new network cases Accelerated DBD is a clear winner, with DBD performing worse than the MIP.

The main reason for the difficulty of the new network cases for the DBD implementation is feasibility of the networks. In the existing network case, the networks are either already feasible or can be made feasible very easily, whereas for the new network case a brand new feasible network must be made from scratch. The standard Benders feasibility cuts used by the DBD implementation are not tight enough to help the algorithm converge, and as such none of the instances are solved to optimality in time, with many failing to find a feasible solution at all. By contrast, the IIS feasibility constraints used in Accelerated DBD allow the problems to be solved significantly faster than the MIP, and solutions were found to previously unsolved instances.

There are three other improvements to the Benders decomposition algorithm we consider here: using analytic Benders cuts (one or two) in the warm start, the budget cover constraints, and using analytic Benders cuts (one or two) during the branch and cut process. We consider the existing and new network cases separately, as the algorithm behaves differently on instances in different cases. After each comparison, we take the best result and keep it for the next comparison.

The main statistic we report is the run time (or achieved optimality gap). We also include the mean and geometric mean of the times for each method across all instances in each case, as well as how many instances each method performed best upon. These help to determine whether or not each method is useful. The first choice we examine is the type of Benders cuts used in the warm

Table 1: Problem sizes for Ghaderi and Jabalameli instances [14]

Inst.	N	L	T	Inst.	N	L	T	Inst.	N	L	T
TP1	20	46	5	TP10	40	162	5	TP19	80	171	5
TP2	20	46	10	TP11	40	162	10	TP20	80	171	10
TP3	20	46	20	TP12	40	162	20	TP21	80	171	20
TP4	20	61	5	TP13	60	180	5	TP22	80	280	5
TP5	20	61	10	TP14	60	180	10	TP23	80	280	10
TP6	20	61	20	TP15	60	180	20	TP24	80	280	20
TP7	40	137	5	TP16	60	205	5	TP25	56	200	5
TP8	40	137	10	TP17	60	205	10	TP26	56	200	10
TP9	40	137	20	TP18	60	205	20	TP27	56	200	20

start.

5.1. Warm start

The warm start is implemented in the same way for all cases: solve the LP relaxation of the master problem, solve the sub-problems for the resulting values of the relaxed variables and check whether or not a Benders optimality or feasibility cut must be added. If so, add the necessary cuts. This is repeated until the objective value converges or no new cuts are added. All feasibility cuts used are the IIS cuts discussed in Section 3.3. There are three options for Benders optimality cuts: non-analytic cuts (NonA) which use the dual variables returned by Gurobi to construct the cut, one analytic cut (AnOne) which is always the λ-heavy cut, or both analytic cuts (AnTwo), which adds both the λ-heavy and γ-heavy cuts if possible. The choice of using the λ-heavy cut over the γ-heavy cut for the AnOne implementation was arbitrary, as the problem can be solved using either type of cut and there was no clear-cut winner from our initial experimentation.

Table 4 contains the results for the existing network case, where it is clear that using both analytic cuts in the warm start is highly beneficial. The mean and geometric means for AnTwo are significantly lower than the other two methods, and it won in 19 of 24 cases. All cases where it was slower are smaller instances which run in under 100 seconds. For many of the larger instances, it makes a significant improvement.

For the new network case, Table 4 shows that using analytic cuts in the warm start becomes a

Table 2: Comparison of runtimes and objective values between the CPLEX implementation in Ghaderi and Jabalameli (2013) and our MIP, standard DBD and accelerated DBD on existing network instances. Runs which reached the time limit are reported as Time (gap%).

Instance	Objective	Ghaderi et. al. (2013) CPLEX				MIP time (s)	DBD time (s)	Accelerated DBD time (s)
		Objective	Lower Bound	Gap (%)	Time (s)			
TP1E	14924.87	14924.87	14924.87	0.00	4	1.24	1.87	0.87
TP2E	29698.23	29698.23	29698.23	0.00	14	3.00	4.26	2.12
TP3E	67676.54	67676.54	67676.54	0.00	81	15.56	12.27	6.05
TP4E	31419.98	31419.98	31419.98	0.00	5	1.62	2.35	1.07
TP5E	67884.16	67884.16	67884.16	0.00	17	5.30	5.80	1.96
TP6E	158091.31	158091.31	158091.31	0.00	77	13.41	13.00	6.00
TP7E	21159.56	21159.56	21159.56	0.00	63	7.24	12.55	8.92
TP8E	45581.23	45581.23	45581.23	0.00	1758	75.48	33.45	24.22
TP9E	111643.23	111643.23	110830.16	0.73	40000	253.72	95.71	62.59
TP10E	19232.89	19232.89	19232.89	0.00	289	44.15	20.82	11.92
TP11E	40310.81	40310.81	39730.11	1.44	20000	664.06	250.15	43.04
TP12E	99032.15	99103.48	96797.34	2.33	40000	1456.10	501.12	95.33
TP13E	25110.09	25110.09	25110.09	0.00	347	31.66	28.67	19.40
TP14E	56892.74	56892.74	56892.74	0.00	1479	89.81	69.40	40.39
TP15E	151049.50	151115.62	149518.50	1.06	60000	2778.78	356.96	324.78
TP16E	26973.29	26973.29	26973.29	0.00	480	36.13	40.98	24.65
TP17E	58636.16	58636.16	58636.16	0.00	16514	273.39	117.04	73.78
TP18E	143986.33	144200.68	142289.95	1.33	60000	972.10	435.64	467.36
TP19E	33696.53	33696.53	33696.53	0.00	5622	74.40	45.38	40.70
TP20E	73969.28	74677.35	73256.40	1.90	40000	481.31	156.72	132.74
TP21E	196674.19	198074.52	192701.54	2.71	80000	16678.98	5884.91	4239.84
TP22E	30999.44	30999.44	30999.44	0.00	1179	59.19	71.87	44.45
TP23E	70070.24	70092.80	69623.72	0.67	40000	622.60	235.13	156.19
TP24E	182700.42	183511.97	179939.71	1.95	80000	9887.72	4536.41	2235.90

Table 3: Comparison of runtimes and optimality gaps between the CPLEX and two heuristic implementations in Ghaderi and Jabalameli (2013) and our MIP, standard DBD and accelerated DBD on new network instances. Runs which reached the time limit are reported as Time (gap%). If no valid solution was found, they are reported as Time (NS).

Instance	Best Objective (Best Bound)	Ghaderi et. al. (2013) Times (s)			MIP Time (s)	DBD Time (s)	Accelerated DBD Time (s)
		CPLEX	Heuristic	Hybrid SA			
TP1N	20473.40	Time (7.02%)	427 (0.00%)	427 (0.00%)	283.05	Time (5.1%)	4.13
TP2N	35528.03	Time (NS)	371 (0.59%)	397 (0.00%)	557.54	Time (1.4%)	11.22
TP3N	70973.67	Time (NS)	604 (1.10%)	645 (0.00%)	8619.04	Time (NS)	28.22
TP4N	35614.49	Time (10.66%)	1289 (2.23%)	1455 (0.00%)	2891.17	Time (9.5%)	5.42
TP5N	69150.51	Time (18.14%)	1208 (1.50%)	1534 (0.62%)	Time (9.2%)	Time (11.2%)	11.00
TP6N	151120.93	Time (NS)	1869 (0.63%)	4497 (0.28%)	Time (6.5%)	Time (13.1%)	25.14
TP7N	23120.55	Time (9.06%)	519 (0.00%)	519 (0.00%)	6500.70	Time (0.6%)	21.06
TP8N	46256.15	Time (NS)	2296 (0.00%)	2296 (0.00%)	Time (2.5%)	Time (NS)	761.75
TP9N	105113.30	Time (NS)	3238 (0.07%)	6501 (0.05%)	Time (2.3%)	Time (NS)	916.12
TP10N	19504.02	Time (NS)	396 (4.63%)	2834 (2.31%)	3801.03	Time (1.6%)	59.80
TP11N	37316.52	Time (8.79%)	1822 (4.03%)	2605 (0.26%)	10577.88	Time (NS)	140.99
TP12N	89331.29	Time (2.68%)	2521 (3.16%)	9828 (1.17%)	10740.89	Time (NS)	1417.29
TP13N	28827.77	Time (9.42%)	7428 (0.14%)	1758 (0.39%)	14074.88	Time (NS)	310.84
TP14N	59320.43	Time (NS)	14627 (2.34%)	26262 (1.28%)	Time (1.2%)	Time (NS)	967.67
TP15N	145023.30	Time (NS)	Time (3.31%)	15819 (1.28%)	Time (0.4%)	Time (NS)	7798.54
TP16N	27741.57	Time (NS)	7782 (2.55%)	8252 (0.55%)	Time (2.6%)	Time (NS)	572.80
TP17N	55740.52	Time (NS)	17136 (2.05%)	21361 (1.19%)	Time (1.5%)	Time (NS)	2587.47
TP18N	132520.16	Time (6.25%)	Time (2.29%)	13660 (1.12%)	39282.31	Time (NS)	2753.99
TP19N	37562.93	Time (19.22%)	10805 (0.71%)	19255 (0.04%)	Time (2.8%)	Time (NS)	1427.63
TP20N	76063.44	Time (NS)	Time (0.22%)	15477 (0.37%)	Time (1.8%)	Time (NS)	13749.00
TP21N	$^{183881.92}_{(183555.16)}$	Time (NS)	Time(2.15%)	24659 (1.16%)	Time (3.5%)	Time (NS)	Time (0.2%)
TP22N	33258.34	Time (6.58%)	8332 (2.06%)	12874 (0.73%)	Time (0.2%)	Time (NS)	454.37
TP23N	70208.93	Time (7.84%)	Time (2.09%)	31868 (1.25%)	Time (0.8%)	Time (NS)	3177.27
TP24N	$^{171442.41}_{(171036.86)}$	Time (NS)	Time (2.48%)	60596 (1.65%)	Time (0.8%)	Time (NS)	Time (0.2%)

Table 4: Comparison of the algorithm's performance when using different cuts in the warm start for existing and new network instances. Non-analytic (NonA), one analytic cut (AnOne) and two analytic cuts (AnTwo). All times are in seconds, instances that are not solved to optimality are marked with Time (optimality gap). Also reported are the Mean time, Geometric Mean (G.M.) time and the number of times each algorithm was best in the instance set.

Instance	Existing network instances			New network instances		
Number	NonA	AnOne	AnTwo	NonA	AnOne	AnTwo
TP1	0.83	0.98	0.70	3.78	5.27	5.69
TP2	2.00	1.94	1.67	9.03	12.13	12.02
TP3	5.84	5.26	4.49	24.76	24.82	23.83
TP4	0.88	0.88	0.97	3.63	8.52	9.47
TP5	1.80	2.06	2.05	7.95	12.92	16.30
TP6	5.25	4.45	4.14	16.76	28.35	29.18
TP7	7.75	7.97	5.83	24.26	42.95	35.22
TP8	20.82	17.03	18.75	290.13	635.40	353.16
TP9	61.00	44.33	35.22	2136.79	1951.36	2664.50
TP10	10.44	13.94	15.19	69.83	398.01	117.59
TP11	112.67	96.11	43.70	145.50	350.51	4912.77
TP12	89.47	136.03	113.27	1361.18	2130.17	1922.54
TP13	19.44	14.09	10.16	295.77	919.66	400.17
TP14	34.27	25.10	21.16	2085.27	735.11	1839.12
TP15	394.73	175.58	154.17	6754.99	29861.33	Time (0.4%)
TP16	22.71	29.31	17.66	568.55	14900.21	699.46
TP17	66.94	63.82	43.07	1952.92	2097.34	3922.23
TP18	337.36	212.04	204.33	2547.31	4818.17	6516.81
TP19	34.16	37.84	25.66	975.16	2227.87	2512.98
TP20	243.60	210.42	108.99	9748.52	Time (1.3%)	15566.78
TP21	4269.29	7600.51	1435.95	Time (0.5%)	Time (8.6%)	Time (1.5%)
TP22	45.06	41.09	18.35	226.32	406.74	1289.39
TP23	184.25	91.47	54.17	2112.00	5488.23	Time (1.08%)
TP24	5875.25	827.90	440.33	Time (0.2%)	Time (0.8%)	Time (0.4%)
Mean time	493.57	402.51	115.83	7965.48	11115.54	12602.53
G.M. time	35.79	30.47	21.66	363.89	713.38	782.02
Wins	3	1	19	21	2	1

hindrance, with non-analytic cuts winning 21 of 24 times. In one instance, NonA was only 0.93s slower than the best method, and the other two instances where it lost, it is not significantly slower. The mean and geometric mean for NonA are significantly lower than those for AnOne and AnTwo.

For the next comparison, the existing network cases will be run using a warm start with two analytic cuts, and the new network cases with a warm start with non-analytic cuts.

5.2. Budget cover constraints

The usefulness of the budget cover constraints is more difficult to determine as they make significantly less difference than the choice of warm start optimality cuts, as can be seen in Table 5. For the existing network case, the means and geometric means of the run times are almost identical, with a roughly 50-50 split on the number of wins. For the new network case, they yield a slight improvement, with lower mean and geometric means of the run times, and winning over 60% of the time.

The difference between the existing network and new network cases is that in the new network case the network must be built from scratch and the budget is twice as big, so there are many more arcs and facilities which will be fractionally open in the warm start solutions. For the existing network case, many arcs and facilities are already open, and so there will be fewer fractional variables in the warm start solution, meaning fewer budget cover constraints will be added.

For the next test, the existing network instances will be run with two analytic cuts and no budget cover constraints in the warm start, and the new network instances with non-analytic cuts and with budget cover constraints.

5.3. Benders optimality cuts

The last choice to make is the type of Benders optimality cuts to use in the main callback. These are the cuts that will be added during the branch and bound process to solve the problem to optimality. Again, the options are the non-analytic dual variables returned by Gurobi (CNonA), or the analytically constructed dual variables. For the one cut case (CAnOne), the λ-heavy cut will be used for the same reason as the warm start.

For the existing network case (seen in Table 6), analytic cuts are clearly beneficial, with CNonA only winning on one of the 24 instances (TP12E). The difference between analytic and non-analytic run times was also clear, with non-analytic having a higher geometric mean than the analytic

Table 5: Comparison of algorithm's performance for existing and new network instances, with and without budget cover constraints.

Existing	Without	With	New	Without	With
TP1E	0.70	0.70	TP1N	3.78	2.55
TP2E	1.67	2.26	TP2N	9.03	7.72
TP3E	4.49	5.38	TP3N	24.76	20.20
TP4E	0.97	0.86	TP4N	3.63	4.66
TP5E	2.05	1.93	TP5N	7.95	10.79
TP6E	4.14	4.63	TP6N	16.76	17.77
TP7E	5.83	5.98	TP7N	24.26	23.55
TP8E	18.75	16.30	TP8N	290.13	203.77
TP9E	35.22	36.05	TP9N	2136.79	1031.04
TP10E	15.19	13.61	TP10N	69.83	56.38
TP11E	43.70	42.65	TP11N	145.50	196.34
TP12E	113.27	97.32	TP12N	1361.18	1180.41
TP13E	10.16	11.16	TP13N	295.77	321.50
TP14E	21.16	25.57	TP14N	2085.27	2311.46
TP15E	154.17	133.59	TP15N	6754.99	8590.36
TP16E	17.66	18.90	TP16N	568.55	515.83
TP17E	43.07	47.96	TP17N	1952.92	1504.40
TP18E	204.33	199.74	TP18N	2547.31	2312.52
TP19E	25.66	26.31	TP19N	975.16	1578.65
TP20E	108.99	73.80	TP20N	9748.52	8750.23
TP21E	1435.95	1610.16	TP21N	Time (0.5%)	Time (0.4%)
TP22E	18.35	21.03	TP22N	226.32	221.79
TP23E	54.17	65.08	TP23N	2112.00	1399.53
TP24E	440.33	383.10	TP24N	Time (0.2%)	Time (0.2%)
Mean	115.83	118.50		7965.48	7919.86
Geometric Mean	21.66	21.89		363.89	343.35
Wins	13	10		9	15

Table 6: Comparison of different Benders optimality cuts for the existing network instances. All implementations use the same code (two-cut analytic warm start and no budget cover constraints) with the exception of the callback used to implement Benders optimality cuts. The different callbacks use non-analytic duals (CNonA), one analytic cut (CAnOne) and two analytic cuts (CAnTwo). All times are in seconds, instances that are not solved to optimality are marked with Time (optimality gap). Also reported are the Mean time, Geometric Mean (G.M.) time and the number of times each algorithm was best in the instance set.

Instance	Existing network instances			New network instances		
Number	CNonA	CAnOne	CAnTwo	CNonA	CAnOne	CAnTwo
TP1	0.70	0.65	0.76	2.55	2.98	2.89
TP2	1.67	1.84	1.39	7.72	7.09	7.86
TP3	4.49	4.32	4.46	20.20	22.43	20.60
TP4	0.97	1.28	0.85	4.66	3.69	3.22
TP5	2.05	2.03	2.16	10.79	8.99	10.54
TP6	4.14	3.86	4.27	17.77	16.11	16.44
TP7	5.83	5.67	5.47	23.55	26.18	22.68
TP8	18.75	14.43	15.00	203.77	236.09	274.26
TP9	35.22	26.65	26.41	1031.04	710.72	1683.92
TP10	15.19	14.26	14.19	56.38	50.86	187.88
TP11	43.70	37.17	37.35	196.34	198.83	197.57
TP12	113.27	155.75	122.36	1180.41	810.04	958.81
TP13	10.16	6.78	6.74	321.50	511.30	467.62
TP14	21.16	15.55	15.27	2311.46	1362.79	987.09
TP15	154.17	128.18	114.79	8590.36	16371.58	16220.59
TP16	17.66	15.22	13.60	515.83	439.66	485.31
TP17	43.07	36.54	34.15	1504.40	1254.36	1315.31
TP18	204.33	140.17	163.00	2312.52	3934.44	2962.72
TP19	25.66	17.20	17.78	1578.65	4851.20	2123.34
TP20	108.99	45.61	49.23	8750.23	25858.42	7123.92
TP21	1435.95	1243.31	1568.76	Time (0.4%)	Time (0.2%)	Time (0.3%)
TP22	18.35	11.73	12.04	221.79	455.68	286.42
TP23	54.17	36.32	31.69	1399.53	973.76	1840.23
TP24	440.33	373.03	515.29	Time(0.2%)	Time (0.2%)	Time (0.2%)
Mean time	115.83	97.40	115.71	7919.86	9080.08	8208.89
G.M. time	21.66	18.10	17.95	343.35	379.61	377.50
Wins	1	12	11	9	10	4

Table 7: Comparison of solution times for problem TP9E with different levels of disaggregation

Disaggregation level	# S.P's solved	S.P. cumulative time (s)	Master solve time (s)
Time only	580	24.74	283.35
Node only	760	14.72	232.77
Node and Time	11200	9.48	101.47

methods. As for whether to use one or two analytic cuts, it seems to make little difference with very similar run times and an even split of the wins. CAnOne won all but one of instances 18-24 (it lost to CAnTwo on TP23E by 5 seconds), however it was often very close.

Table 6 shows that for the new network case, using analytic cuts was again less beneficial, with CAnOne having a lower mean and geometric mean of the run times. CAnTwo only won a few times, with the rest of the instances being split between CNonA and CAnOne. For instances 18 and above, CNonA beat CAnOne in all but two cases (TP21N and TP23N) and tied in one (TP24N).

6. Discussion

6.1. Disaggregation level

In this problem it is possible to disaggregate over two different sets: the source of each demand and each separate time period. We show here that it is best to disaggregate by both sets at the same time. Disaggregation of sub-problems, and thus Benders cuts, always results in tighter bounds. These tighter bounds allow the master problem to be solved more quickly. The trade-off is that having more sub-problems can take longer to solve, particularly if there are overheads associated with those sub-problems. In this problem, the most sub-problems we solve are 1600, which is acceptable considering the speed increases we obtain from this. In other problems, the number of sub-problems may enter the hundreds of thousands, at which point even the smallest overheads will start to add up.

A specific example is data set TP9E, which we can compare results for if we disaggregate only by nodes, only by time and by both nodes and time. For this instance there are 40 nodes and 20 time periods. Table 7 shows the number of sub-problems (S.P.'s) solved, the total time spent solving sub-problems and the total time spent solving the entire problem for the different levels of disaggregation.

We can see that disaggregating more leads to smaller sub-problems which solve significantly faster. The average solve time for each sub-problem is 43ms, 19ms and 0.85ms for time only, nodes only and both, respectively. Even though many more sub-problems must be solved when disaggregating by both nodes and time, the cumulative time spent solving them is less, and the tighter cuts provided by disaggregation leads to a faster solve time of the master problem.

7. Conclusion

Disaggregated Benders decomposition with lazy constraints is an effective method for solving the DUFLNDP if implemented properly. Adding constraints that enforce feasibility to avoid relying upon Benders feasibility cuts, and using a warm start are good ways of improving the effectiveness of the solver. Analytically derived Pareto-optimal Benders cuts can also be beneficial in some cases. For this particular problem, it is the disaggregation of the sub-problems and IIS feasibility cuts which provide the impressive speed increase, which has allowed us to solve almost all instances to optimality within the time limits. In the future, we would like to generalise this approach to a wide range of network design and facility location problems where similar techniques are beneficial.

8. References

References

[1] A. Weber, Uber den Standort der Industrian, University of Chicago Press, 1909.

[2] J. Benders, Partitioning procedures for solving mixed-variables programming problems, Numerische Mathematik 4 (1962) 238–252.

[3] M. Fischetti, I. Ljubic, M. Sinnl, Redesigning benders decomposition for large-scale facility location, Management Science.

[4] A. Geoffrion, G. Graves, Multicommodity Distribution System Design by Benders Decomposition, International Series in Operations Research and Management Science, Springer Science+Business Media, 1974, Ch. 4, pp. 35–61.

[5] T. Magnanti, R. Wong, Accelerating benders decomposition: Algorithmic enhancement and model selection criteria, Operations Research 29 (3) (1981) 464–484.

[6] L. Tang, W. Jiang, G. Saharidis, An improved benders decomposition algorithm for the logistics facility location problem with capacity expansions, Annals of Operations Research 210 (2013) 165–190. doi:10.1007/s10479-011-1050-9.

[7] L. Ford Jr., D. Fulkerson, Maximal flow through a network, Canadian Journal of Mathematics 8 (1956) 399–404.

[8] A. F. P. Elias, C. Shannon, A note on the maximum flow through a network, IRE Transactions on Information Theory IT-2 (1956) 117–119.

[9] S. Nurre, B. Cavdaroglu, J. Mitchell, T. Sharkey, W. Wallace, Restoring infrastructure sytems: An integrated network design and scheduling (INDS) problem, European Journal of Operational Research 223 (2012) 794–806. doi:10.1016/j.ejor.2012.07.010.

[10] N. Boland, T. Kalinowski, H. Waterer, L. Zheng, Scheduling arc maintenance jobs in a network to maximize total flow over time, Discrete Applied Mathematics 163 (2014) 34–52. doi:10.1016/j.dam.2012.05.027.

[11] E. Martins de Sá, I. Contreras, J. Cordeau, R. de Camargo, G. de Miranda, The hub line location problem, Transportation Science 49 (3) (2015) 500–518. doi:10.1287/trsc.2014.0576.

[12] R. de Camargo, G. Miranda Jr., H. Luna, Benders decomposition for the uncapacitated multiple allocation hub location problem, Computers and Operations Research 35 (2008) 1047–1064. doi:10.1016/j.cor.2006.07.002.

[13] E. de Sá, R. de Camargo, G. de Miranda, An improved benders decomposition algorithm for the tree of hubs location problem, European Journal of Operational Research 226 (2013) 185–202. doi:10.1016/j.ejor.2012.10.051.

[14] A. Ghaderi, M. Jabalameli, Modeling the budget-constrained dynamic uncapacitated facility location-network design problem and solving it via two efficient heuristics: A case study of health care, Mathematical and Computer Modelling 57 (2013) 382–400. doi:10.1016/j.mcm.2012.06.017.

[15] Gurobi Optimization, Inc., Gurobi optimizer reference manual (2015).
URL http://www.gurobi.com

[16] Z. Gu, G. Nemhauser, M. Savelsbergh, Lifted cover inequalities for 0-1 integer programs: Computation, INFORMS Journal on Computing 10 (4) (1998) 427–437.

[17] D. McDaniel, M. Devine, A modified benders' partitioning algorithm for mixed integer programming, Management Science 24 (3) (1977) 312–319.